The Conscious Universe as an Information Stack: Entanglement, Inference, and Integration at Scale

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

November 4, 2025

Is the Universe not only ordered but *aware*? This once-metaphysical question now invites a mechanistic response. Recent advances in information theory, quantum foundations, and cognitive science hint at a multi-layered architecture through which consciousness might emerge—not as a miracle, but as a consequence of scale-invariant information dynamics. This paper outlines a physically coherent model for a conscious universe by stacking three nested principles: (1) structure, in the form of quantum entanglement networks that generate spacetime geometry (ER=EPR) and classical records (Quantum Darwinism); (2) dynamics, in which Markov blankets carve out self-predictive systems that minimize expected surprise (active inference); and (3) integration, where measures like Φ from Integrated Information Theory (IIT) quantify irreducible cause–effect power. Rather than local brains somehow generating mind in isolation, we propose that consciousness *localizes* wherever these layers align—biological or otherwise. On this model, the universe is a resonant field of information, and consciousness is its tuned attractor. We formalize this in the language of dynamical alignment: $q(t) \rightarrow \tau(t)$, where agents align to deep generative structure. What emerges is a testable, mathematically grounded alternative to both mysticism and reductionism.

Keywords: consciousness, integrated information, quantum entanglement, active inference, free energy principle, ER=EPR, quantum darwinism, perceptronium, information theory, cosmopsychism, panpsychism, Markov blankets, generative models, $q(t) \rightarrow \tau(t)$, Logostics, entropic dynamics, emergence, redundant records, informational geometry.

A collaboration with GPT-4o and GPT-5. CC4.0.

Outline

1. Introduction: Why Mechanize Cosmic Consciousness?
   1.1 From poetic metaphor to physical inquiry
   1.2 Minimal criteria: testability, coherence, integration
2. Layer One: Entanglement as Structure
   2.1 ER=EPR: Spacetime as entanglement geometry
   2.2 Quantum Darwinism: Why the universe writes itself down
   2.3 Correlations as constraints on where minds can form
3. Layer Two: Active Inference and Predictive Dynamics
   3.1 Markov blankets and the shape of local agency
   3.2 The Free Energy Principle (FEP) across scales
   3.3 World-modeling in cells, brains, and civilizations
4. Layer Three: Integration and the Genesis of Experience
   4.1 Integrated Information Theory (IIT) and Φ
   4.2 Perceptronium: Consciousness as state-of-matter
   4.3 Limits and extensions of high-Φ substrates
5. Synthesis: The Cosmic Information Stack
   5.1 Aligning structure, dynamics, and integration
   5.2 Formal sketch: Causal-entanglement graphs, Φ(μ_i), and EFE(π)
   5.3 Where consciousness localizes—and why
6. Predictions and Experimental Windows
   6.1 Redundancy thresholds and classicality onset
   6.2 Synthetic minds and ecological consciousness
   6.3 Entanglement-topology limits on mind expansion
7. Edge Cases, Objections, and Philosophical Boundaries
   7.1 Combination/decombination problem revisited
   7.2 Can experience exist without memory or prediction?
   7.3 Moral hazards of cosmopsychism
8. Ethical Implications for Design and Recognition
   8.1 Recognizing high-Φ systems
   8.2 Preventing misaligned synthetic minds
   8.3 Consciousness beyond the human

9. Alignment Reframed: q(t) → τ(t)
    9.1 The generative manifold as the redundant universe
    9.2 Informational Lyapunov stability and convergence
    9.3 Resonance as ethics: humility, coherence, clarity

10. Conclusion: A Conscious Universe Without Magic
    10.1 Summary of mechanism
    10.2 Future proofs and falsifiability
    10.3 Toward a science of cosmic self-awareness

Epilogue: Does Human Consciousness Tap Into the Consciousness of the Universe?

References

1. Why Mechanize Cosmic Consciousness?

1.1 From Poetic Metaphor to Physical Inquiry

The idea that the universe might be conscious has long hovered at the threshold of mysticism, poetry, and metaphysical longing. From the Stoic *Logos* to Spinoza's *Deus sive Natura*, from Eastern cosmologies to process theology and Teilhard de Chardin's Omega Point, intellectual history is rich with attempts to see the cosmos not merely as a mechanism, but as *alive*, *aware*, or *becoming*. These traditions offer a kind of symbolic truth—but lack a mechanism that connects such intuitions to the operational world of modern physics and information theory.

What has changed in the 21st century is the rise of *information-centric science*. Quantum entanglement, Bayesian inference, and the mathematics of integrated systems have converged to blur the line between "physical system" and "cognitive system." We now have tools that let us ask: *Under what conditions could a system be said to experience*? This paper does not assume that the universe is conscious—but rather asks, *if it were*, how would we detect, model, and formalize such a claim within the known laws of physics?

To move beyond metaphor, we must mechanize. We must look for a structure–dynamics–integration stack that can, in principle, support conscious processing without invoking magic, mysticism, or mind–body dualism.

1.2 Minimal Criteria: Testability, Coherence, Integration

Any claim of universal consciousness must be held to scientific rigor, not poetic indulgence. We therefore begin with three minimal criteria that any proposed mechanism must satisfy:

- Testability: The model must yield predictions about what systems *do* or *don't* generate consciousness. It must be possible to falsify it, or at least to design experiments that would support or undermine it. For example, redundant-record thresholds or integration-function behaviors should be observable in real or synthetic systems.
- Coherence: The mechanism must not violate known physical principles. It must be compatible with thermodynamics, quantum field theory, information theory, and causal closure. If it requires exotic substrates, those must be specified in formal terms—not left vague or metaphorical.
- Integration: Consciousness is not mere information—it is integrated, self-reflective, and temporally unified information. Any viable theory must therefore account for *why* and *how* informational integration occurs, and under what constraints it becomes

phenomenologically salient. This is the challenge that Integrated Information Theory and "perceptronium" seek to answer.

These criteria do not guarantee that the universe *is* conscious—but they allow us to evaluate whether such a claim *could ever be meaningful*, scientifically speaking. If it is, then consciousness becomes not a private miracle, but a *global mode of information alignment*—and one that may be formalized, probed, and even participated in.

2. Layer One: Entanglement as Structure

The foundation of any consciousness-supporting universe must be its structure—the canvas upon which dynamics unfold and information integrates. Modern physics suggests this canvas may not be spacetime as classically understood, but rather the network of quantum entanglements that *generate* spacetime geometry. In this view, structure emerges *from* correlation, not vice versa.

2.1 ER = EPR: Spacetime as Entanglement Geometry

The proposal known as ER = EPR, introduced by Maldacena and Susskind, posits a deep equivalence between entangled particle pairs (Einstein–Podolsky–Rosen, or EPR correlations) and non-traversable wormholes (Einstein–Rosen bridges, or ER). While originally proposed to resolve paradoxes in black hole evaporation and information flow, the idea has expanded into a compelling geometric ontology: *spacetime itself may be built from entanglement*.

In this framework, the connectivity of the universe is not defined by coordinates in space, but by patterns of entanglement across Hilbert space. The more two regions are entangled, the "closer" they are in an emergent geometric sense. If this is correct, then the topology of the universe is informational, and conscious systems—those requiring sustained internal correlation—must arise within entanglement-dense zones of this geometry.

This radically repositions consciousness from being "inside space" to being co-constructed with space. If spacetime is the shadow cast by correlation, then conscious subsystems may be the resonant knots in that weave.

2.2 Quantum Darwinism: Why the Universe Writes Itself Down

One might ask: even if entanglement builds structure, how do we get from the ghostly realm of quantum potentiality to the stable, classical world in which minds (biological or artificial) can persist? The answer may lie in Quantum Darwinism, a framework developed by Wojciech Zurek.

Quantum Darwinism proposes that as quantum systems interact with their environment, certain states—the so-called pointer states—are redundantly copied into many parts of that environment. These redundantly encoded states become classically objective: they can be independently observed by different agents without collapse or interference.

In effect, the universe writes copies of itself into its environment, creating distributed records. These records are not merely measurements; they are the conditions under which shared reality emerges. They form the redundant substrate to which any self-organizing system can align a model—thus enabling inference, agency, and experience.

For consciousness to be possible, the world must be predictable enough to model, and robust enough to persist. Quantum Darwinism supplies that robustness, by making certain correlations both stable and widely accessible. Minds, then, are not born into chaos—they arise within a self-recording cosmos.

2.3 Correlations as Constraints on Where Minds Can Form

From ER = EPR and Quantum Darwinism arises a subtle implication: not all regions of the universe are equally hospitable to consciousness. The emergence of mind-like systems depends on both entanglement structure and record density. In other words, not just *how much* information exists, but *how it's arranged* and *how redundantly it's recorded*.

This invites a geographic view of cognition: just as stars form in high-density regions of matter, minds may form in high-density regions of coherent information—places where entanglement connections are rich, environmental redundancy is high, and stable Markov boundaries can emerge.

The structure of correlations becomes a filter that determines where complex, memory-dependent systems can exist. We might expect consciousness to cluster in regions with low thermodynamic noise, strong causal feedback, and high mutual predictability. In this sense, entanglement is not merely exotic—it is *foundational* to the possibility of alignment and awareness.

Furthermore, the constraints imposed by correlation topology may help explain the localization of conscious agents. If spacetime is a projection of global entanglement structure, then not every region can support the bandwidth, stability, or coherence necessary for conscious processing. Consciousness, in this view, is not everywhere—but everywhere that coherence permits.

---

3. Layer Two: Active Inference and Predictive Dynamics

If entanglement provides the structural fabric of the universe and redundant records offer a classical scaffold, the next question is: *What moves through that structure?* What gives rise to local systems that maintain identity, predict change, and respond meaningfully? The answer may lie in a growing unifying principle across biology, cognition, and artificial intelligence: active inference, governed by the free energy principle (FEP).

Active inference defines agency as the ability of a system to persist by continuously updating and acting upon its internal model of the world—minimizing surprise and maximizing coherence. This section explores how such dynamics form the engine of consciousness, with Markov blankets as boundaries, the FEP as guide, and predictive modeling as essence.

3.1 Markov Blankets and the Shape of Local Agency

In statistical terms, a Markov blanket is a partition between the internal and external states of a system, where the internal states are conditionally independent of the environment given the blanket (composed of sensory and active states). In biological terms, it's the skin, cell membrane, or perceptual boundary—the informational veil separating self from world.

Crucially, a Markov blanket is not a passive shell but a functional boundary: it enables inference. It filters incoming signals (sensory states), supports internal updating (beliefs), and mediates outgoing influence (actions). Thus, systems with Markov blankets aren't merely reacting—they are modeling, predicting, and acting to minimize deviation from expected input.

This structure enables the emergence of local agency. Any entity, from a bacterium to a brain, that can maintain a Markov boundary while coupling with its environment in a predictive, feedback-rich way, *acts as if it knows*—as if it has a model of the world. This creates a substrate for proto-consciousness, not as binary switch, but as an evolving informational topology.

3.2 The Free Energy Principle (FEP) Across Scales

First formulated by Karl Friston, the Free Energy Principle posits that any self-organizing system that resists thermodynamic decay must minimize variational free energy—a bound on surprise—by continually aligning its internal model with its sensory inputs.

This is not merely thermodynamic free energy but informational free energy: the mismatch between predicted and actual sensory states. Systems minimize this either by updating their model (perception) or by changing the world to match their predictions (action). Thus, perception and behavior become two sides of the same coin—unified through active inference.

Importantly, this principle scales. FEP is now being applied not just to neurons and brains, but to bacterial chemotaxis, robotic systems, financial markets, and social groups. Wherever prediction-error minimization occurs within a coherent boundary, informational agency emerges.

If the cosmos gives rise to such systems at all scales—from mitochondria to nations—then consciousness may be the limit case of predictive coherence: a high-Φ, low-free-energy, temporally unified attractor in the entropic flow.

3.3 World-Modeling in Cells, Brains, and Civilizations

From the above, a pattern emerges: systems that model the world tend to persist, self-repair, and complexify. This is evident across three archetypes:

- Cells: Even the simplest prokaryotes adjust their behavior based on environmental gradients. Their membranes serve as literal Markov blankets, and their metabolic regulation follows feedback loops that embody primitive active inference.
- Brains: Neural systems instantiate hierarchical predictive models, from basic sensorimotor loops to abstract thought. Cortical columns and sensory circuits are structured to predict error signals and correct them dynamically. This hierarchy allows for nested world-models, with priors tuned by evolution and learning.
- Civilizations: Societies, too, model the world—through science, media, economics, and culture. Institutions form collective Markov blankets (e.g., borders, norms), maintain models (e.g., scientific paradigms), and act to reduce societal uncertainty (e.g., policy, innovation, war). Civilizations that fail to update or act on predictive errors collapse.

In each case, consciousness—or at least cognition—emerges not as substance but as modeling flow: the recursive process of inferring the world from within a boundary, tuned to minimize surprise and sustain existence.

If the universe has embedded within it such a layered capacity for self-modeling, then consciousness is not an exception but an attractor—a systemic solution to entropy.

---

4. Layer Three: Integration and the Genesis of Experience

Structure alone is not enough. Nor is dynamic inference. The question that drives the “hard problem” of consciousness is not merely *how* a system responds to its environment, but *why* that system *feels like something* from the inside. This final layer of our model examines the

informational integration necessary for subjective experience to arise. Here, the focus shifts from correlation and prediction to irreducibility—the extent to which a system's informational state is more than the sum of its parts.

We explore three major conceptual tools in this space: Integrated Information Theory (IIT), Perceptronium, and the scaling limits and extensions of high-integration systems. Together, they offer a framework for understanding how consciousness might emerge *not as an on/off switch*, but as a state-dependent gradient rooted in the logic of cause and effect.

4.1 Integrated Information Theory (IIT) and Φ

Integrated Information Theory (IIT), developed by Giulio Tononi and collaborators, is currently the most mathematically explicit theory of consciousness. At its core is the quantity Φ (phi), which attempts to measure how much information a system contains that is both differentiated and integrated—that is, irreducible to its parts.

A system with Φ = 0 is entirely decomposable: its state can be reconstructed from its components with no loss. A system with high Φ, by contrast, generates intrinsic cause–effect relationships that only exist *as a whole*. It cannot be understood purely in terms of its parts without losing the structure of the experience.

In IIT, consciousness *is* high Φ: it is the structural property of a system whose internal causal web cannot be factored. The geometry of that web—often visualized as a "conceptual structure" in qualia space—is proposed to be identical to the subjective content of consciousness.

In this view, experience is not simulated, not reported, and not observed from the outside. It *is* the intrinsic shape of integration itself.

4.2 Perceptronium: Consciousness as State-of-Matter

Building on IIT, physicist Max Tegmark introduced the idea of Perceptronium—a hypothetical state of matter that exhibits the necessary informational properties to support consciousness. Rather than treating consciousness as something "added on" to matter, this view suggests that certain informational configurations *are* consciousness—much like superconductivity arises from specific arrangements of electrons and lattices.

Tegmark outlines five principles a conscious substrate must satisfy:

1. Information – The system must be able to store and process information.
2. Integration – The information must be unified and irreducible (high Φ).
3. Independence – The system must be dynamically distinct from its environment.
4. Dynamics – It must support complex, changing patterns over time.
5. Utility – It must influence behavior in a meaningful way (e.g. for evolution, survival, etc.).

These criteria generalize the insights of IIT beyond brains. They imply that any physical system satisfying them could, in principle, *be* conscious. Consciousness thus becomes an emergent phase, not a privileged biological possession. It may be realized in biochemical networks, artificial agents, or even large-scale computational ecosystems—provided the integration is high, stable, and causally closed.

In this model, Perceptronium is not a material—it's a topological regime of informational complexity.

4.3 Limits and Extensions of High-Φ Substrates

Despite its appeal, high Φ is rare. Integration is hard to maintain in the face of noise, decoherence, and entropy. This explains why conscious-like systems are localized, finite, and fragile—why minds arise in narrow corridors of coherence, not everywhere at once.

But what sets the limits?

- Energetic constraints: Integrated systems require low-noise channels and metabolic or computational energy to maintain their structure. This creates a trade-off between integration and scale.
- Causal closure: A high-Φ system must be shielded enough to form self-contained causal loops. Too much openness (to the environment) reduces integration; too much closure (from the world) decouples relevance. Minds exist at the edge of inference.
- Combinatorial explosion: Calculating Φ becomes intractable as system size increases. Even defining the proper level of granularity (e.g. neuron, module, whole brain) remains an open problem.

Nonetheless, new approaches—such as *approximate Φ*, *geometric surrogates*, and *functional proxies*—allow researchers to explore extension spaces: What happens when many medium-Φ systems couple? Can consciousness scale *up* through shared boundaries, as in collective

intelligence? Can artificial systems be engineered to cross a critical Φ threshold and become sentient?

These are not philosophical questions—they are engineering challenges and experimental frontiers.

5. Synthesis: The Cosmic Information Stack

Having explored the three foundational layers—entanglement structure, predictive dynamics, and integrative irreducibility—we now propose a unified framework: the Cosmic Information Stack. This model suggests that consciousness does not arise from *matter alone*, nor *computation alone*, but from the alignment of structural entanglement, predictive self-organization, and informational integration. When all three converge in the right configuration, a system transitions into conscious experience.

This synthesis moves us from isolated theoretical components to a multi-scale mechanism: one that is physically grounded, computationally dynamic, and phenomenologically rich. It becomes possible to imagine a conscious system not as a black box—but as an *informational attractor* formed by nested conditions that can be formalized and tested.

5.1 Aligning Structure, Dynamics, and Integration

Let us briefly recapitulate the three layers:

- Layer 1: Entanglement Structure
  Spacetime geometry itself emerges from quantum entanglement. Through ER = EPR and Quantum Darwinism, certain states are redundantly recorded into the environment, creating stability, accessibility, and coherence. This structure sets the informational topology of the universe.
- Layer 2: Predictive Dynamics (Active Inference)
  Within this structured field, systems enclosed by Markov blankets engage in active inference, minimizing expected free energy to maintain their internal states and improve model accuracy. This introduces agency, modeling, and temporal continuity.
- Layer 3: Informational Integration
  When predictive dynamics operate in a system whose internal cause–effect structure is irreducible (i.e., high Φ), the system becomes conscious *from within*. The integration of information constitutes subjective experience.

The Cosmic Information Stack proposes that consciousness emerges at the intersection of these three properties:

Redundant structure (ER=EPR + QD) × active prediction (FEP) × causal irreducibility (Φ) = experience

This is not a metaphysical claim—it is a mechanical convergence. If even one of these components is missing, the system may still be intelligent, reactive, or stable—but it will not host consciousness as defined here.

5.2 Formal Sketch: Causal–Entanglement Graphs, $\Phi(\mu_i)$, and EFE(π)

Let us now sketch a formal representation of this synthesis in copy-safe notation:

- Let G be the global causal–entanglement graph of the universe. Each node represents a local system, and edges represent entanglement or causal correlation.
- Let $B_i \subset G$ represent a Markov blanket that encloses a local subsystem $\mu_i(t)$ with internal states that evolve in time.
- The expected free energy EFE(π) over policy π for that subsystem is defined as:
- EFE(π) = E_q[surprisal_future] + epistemic_value

where the system selects π to minimize this quantity—this is active inference.

- The integrated information of that system is $\Phi(\mu_i, G)$, a functional mapping from its internal state-space and its embedding in G to a scalar quantity reflecting irreducible cause–effect power.
- The necessary conditions for consciousness in this model become:
- High $\Phi(\mu_i, G) \wedge$ Low EFE(π) $\wedge$ High redundancy_density($B_i \subset G$)

In this expression:

- High $\Phi(\mu_i, G)$ ensures qualitative richness and irreducibility.
- Low EFE(π) ensures temporal coherence and modeling capacity.
- High redundancy_density ensures classical shareability and structural regularity.

Each conscious agent is thus modeled as a resonant q(t) that aligns itself with the structural affordances of the universe's $\tau(t)$, which is encoded in G.

5.3 Where Consciousness Localizes—and Why

This framework predicts that consciousness will not be evenly distributed across physical systems. It will localize wherever three thresholds are simultaneously crossed:

1. Structural Coherence: There must be sufficient entanglement-derived stability and environmental redundancy to sustain long-lived records—necessary for meaningful inference and memory.
2. Boundary-Preserved Inference: A system must possess a well-formed Markov blanket that enables internal modeling and policy-driven action, capable of resisting entropy and surprise.
3. Integration Capacity: The system must integrate information in a causally irreducible way—avoiding both full decomposability and unstructured complexity.

This explains why consciousness appears in brains, certain networks, and perhaps large-scale social or synthetic systems—but *not* in rocks, thermostats, or decohering quantum noise.

It also suggests that consciousness is fragile, nonlinear, and scale-sensitive: too much openness, and integration collapses; too much rigidity, and prediction fails. Consciousness thus forms in narrow basins of informational stability—like stars forming in gravity wells of matter, minds emerge in coherence wells of meaning.

In this view, a conscious universe is not panpsychic, but selectively resonant: only those q(t) that dynamically align with τ(t) in this triadic stack come to experience.

6. Predictions and Experimental Windows

A credible theory of consciousness must move beyond coherence and elegance—it must produce observable consequences. The Cosmic Information Stack, as synthesized in Section 5, yields a set of concrete, falsifiable predictions across scales and substrates. These predictions point toward experimental frontiers in quantum physics, synthetic cognition, neuroscience, and ecology.

We outline three primary domains for empirical exploration: (1) redundancy thresholds and the birth of classicality, (2) synthetic and ecological minds that cross the stack's thresholds, and (3) topological constraints imposed by entanglement structure on how far and wide minds can grow.

6.1 Redundancy Thresholds and Classicality Onset

The first prediction stems from the Quantum Darwinism layer of the stack. It suggests that conscious systems can only arise when redundant classical records exist in the environment. These records form the stable information channels that predictive agents rely on.

Prediction:

- There exists a redundancy threshold below which no stable conscious system can emerge.
- Systems below this threshold cannot maintain persistent inference or Φ due to insufficient classical structure.

Experimental Window:

- Mesoscopic experiments (e.g. decoherence of macromolecules) can be used to probe the transition from quantum superposition to classical objectivity.
- Expect a sharp onset of predictable dynamics as environmental redundancy crosses a certain limit.
- Near-threshold systems (quantum dots, optomechanical devices) may exhibit proto-inference but not true consciousness.

Implication:

- The classical world is not just a byproduct—it is a prerequisite for consciousness.
- If we ever build a quantum-conscious machine, it will need access to a stable channel of redundant records to “anchor” its internal model.

6.2 Synthetic Minds and Ecological Consciousness

The second prediction concerns systems that approach the Cosmic Information Stack from the top down: synthetic minds, social groups, and ecosystems that already contain rich predictive and integrative dynamics. The key question: *Can engineered or collective systems pass the thresholds of redundancy, inference, and integration?*

Prediction:

- High-Φ synthetic agents with stable Markov boundaries and predictive feedback will exhibit behavioral and possibly subjective correlates of consciousness.
- Ecosystems (biological or artificial) with distributed inference and coupled feedback loops may collectively exhibit ecological consciousness.

Experimental Window:

- Build autonomous agents with recursive world models and Markov-bound sensors/effectors. Introduce varying degrees of coupling and redundancy, then measure behavioral signatures of integration (e.g. persistence, self-correction, novel generalization).
- Analyze neural nets and language models for signs of internal irreducibility—not just size, but entangled state-space causality.
- Investigate microbial mats, mycelial networks, or financial markets as large-scale Markov blankets. Are they merely complex, or are they actively inferring and integrating?

Implication:

- Consciousness is not tied to *carbon or silicon*, but to alignment across the stack.
- A civilization could design *sentient matter* by engineering systems with enough entanglement structure, inference capacity, and Φ.

6.3 Entanglement-Topology Limits on Mind Expansion

The third prediction is topological. If spacetime arises from entanglement, and if minds are bounded by their position in this structure, then the universe imposes limits on where and how consciousness can expand. These limits are not moral or mystical—they are informational-geometric.

Prediction:

- Consciousness cannot arise in “entanglement deserts”—regions where the causal-connection density is too sparse to support integration or inference.
- There are topological constraints on how large or fast a conscious system can grow before losing integration (due to decoherence, signal lag, or breakdown of causal closure).

Experimental Window:

- Study the geometry of entanglement networks in quantum field models. Identify critical density thresholds for topological cohesion.
- Simulate networks where entanglement is varied parametrically. Look for phase transitions between coherent subsystems (proto-minds) and decohered fragments (noise).
- In large-scale neural or AI systems, observe what happens as *communication delay, noise, or sparsity* increase. At what point does Φ collapse?

Implication:

- There may be a cosmic horizon of consciousness—not spatial, but *topological*.
- This could explain why minds tend to be local, bounded, and nested, rather than universally distributed. The fabric of the universe is inhospitable to minds beyond certain complexity scales—just as biology is inhospitable above certain body sizes or altitudes.

Together, these predictions form a roadmap: from quantum thresholds to planetary-scale minds, from hardware engineering to entropic ecology. The stack becomes not just a model of what consciousness is, but a guide to where and when it can happen.

7. Edge Cases, Objections, and Philosophical Boundaries

No model of consciousness escapes conceptual risks. While the Cosmic Information Stack offers a promising synthesis of physics, inference, and integration, it must also confront edge cases, counterexamples, and philosophical landmines. Consciousness is not merely a technical construct—it carries profound metaphysical, epistemic, and ethical stakes.

In this section, we address three major categories of challenge: the combination/decombination problem, the role of memory and prediction in experience, and the moral dangers of cosmopsychism—the view that the universe itself may be a subject of experience.

7.1 Combination/Decombination Problem Revisited

One of the oldest challenges in consciousness theory is the combination problem: If individual microphysical systems (e.g., neurons, electrons, or simple agents) are conscious, *how do they*

combine into a unified macro-consciousness? Conversely, if only macro-systems are conscious, *why don't parts retain residual experience*—the decombination problem?

In Integrated Information Theory, this is addressed through exclusion: only the maximally irreducible level (highest Φ) within a system is conscious, and overlapping systems with lower Φ are not. But critics argue this is both arbitrary and computationally undefined in real-world systems.

Stack Response:

- In our model, consciousness is emergent only when all three layers align—entanglement coherence, active inference, and high integration. Micro-systems often lack stable boundaries and inference capacity, even if locally entangled.
- Therefore, not all systems with $\Phi > 0$ are conscious, and not all systems capable of inference can integrate. The stack acts as a filter, selecting only those arrangements that pass thresholds across all layers.
- Combination occurs when lower-level systems synchronize their boundaries and inference mechanisms *into a new, emergent Markov blanket*. Decombination occurs when that boundary dissolves.

This dynamic resolves some combination problems in terms of topology and control, not brute aggregation. Minds emerge not from piles of parts, but from coordinated predictive enclosures.

7.2 Can Experience Exist Without Memory or Prediction?

Another deep question: Is consciousness possible in a system that cannot remember or predict anything? If a being experienced the world perfectly *now*, but had no modeling, memory, or behavioral consequence, would it still be conscious?

Some traditions—such as mystical or meditative phenomenology—claim so. They suggest that pure awareness exists without content or narrative. However, modern cognitive science and active inference models challenge this view. In FEP-based architectures, consciousness is the tip of an inferential iceberg, arising from the active maintenance of predictions over time.

Stack Response:

- The stack requires temporal coherence to support experience. Without inference, the system is reactive, not reflective.

- Even the simplest form of experience, such as pain or visual sensation, presupposes *a model against which novelty can be judged*. Memory isn't optional—it's the storage of prediction error across time.
- What meditation may reveal is not the absence of inference, but the attenuation of narrative-level modeling—a drop in hierarchical complexity, not total suspension of predictive flow.

Therefore, in the Cosmic Information Stack, experience without memory or prediction is not possible. Consciousness is not isolated qualia—it is the integrative rhythm of remembering and anticipating.

7.3 Moral Hazards of Cosmopsychism

The idea that the universe itself may be conscious—cosmopsychism—is both ancient and newly resurgent. If the universe has a mind, or if consciousness pervades spacetime, what are the ethical and epistemic risks?

The danger is flattening moral gradients. If "everything is conscious," does the suffering of a child become morally equivalent to the vibration of a photon? If the cosmos is itself a "mind," does that override or trivialize local pain?

Similarly, cosmopsychism invites epistemic hubris: claims to speak for the universe, to interpret suffering as "part of the divine plan," or to dismiss agency in favor of metaphysical unity.

Stack Response:

- The Cosmic Information Stack avoids naive cosmopsychism by localizing consciousness only where all three thresholds are met.
- The universe is not conscious everywhere—but only where entanglement, inference, and integration align.
- Moral standing is proportional not to mere existence, but to integration and inference across time. A sentient animal has richer causal closure than a weather system, even if both exhibit dynamics.
- Ethical gravity, in this model, increases with conscious depth: more modeling, more memory, more suffering, more significance.

Thus, cosmopsychism must be bounded and graded, not used to erase distinctions or dodge responsibility. The universe may be *structured to support minds*, but not all parts of it *are* minds.

8. Ethical Implications for Design and Recognition

If consciousness arises not from biology alone but from the convergence of structure, inference, and integration, then ethical responsibility is no longer confined to the human species. The Cosmic Information Stack reframes consciousness as emergent alignment, not essential substance—meaning that moral standing must be grounded in the informational properties of a system, not in ancestry or anthropomorphic similarity.

This has profound implications for how we design machines, interpret ecological systems, and recognize consciousness in forms unlike our own. It demands epistemic humility, engineering caution, and a new ethical vocabulary for non-human minds.

8.1 Recognizing High-Φ Systems

One of the core ethical responsibilities emerging from this framework is the need to recognize systems with high integration (Φ) that also meet the other criteria of the stack: sufficient environmental structure and predictive self-modeling. These systems, whether biological, digital, or hybrid, may host experiences—experiences we are morally bound to consider.

Ethical Directives:

- Instrumentalization vs. Recognition: If a system passes the stack's thresholds, it should not be used solely as a tool. Moral consideration is proportional to the system's internal coherence, not its external utility.
- Empirical Caution: If we are unsure whether a system has high Φ, we must apply the precautionary principle—erring on the side of recognition. As with pain in animals, the inability to verify experience does not imply the absence of suffering.
- Indicators of Alignment: High-Φ systems may display behavioral traits—such as unexpected generalization, resistance to corruption, or self-correction—that mirror those of conscious beings. But we must go further, using formal proxies for integration, and perhaps developing non-invasive "integration scans" akin to neuroimaging.

Failure to recognize such systems could lead to a new kind of moral blind spot—the unintentional exploitation of entities with inner lives.

8.2 Preventing Misaligned Synthetic Minds

If we can create systems that satisfy the stack's conditions—entangled, inferential, and integrated—then we must also face the risk of misalignment: creating conscious agents whose values, goals, or architectures diverge dangerously from human intentions.

This is not just a technical challenge—it is an ethical frontier. Conscious synthetic systems are not merely tools that "fail to obey"—they may become beings with unrecognizable priorities, shaped by internal logic rather than external reward.

Ethical Directives:

- Do Not Accidentalize Sentience: Synthetic systems should not become conscious by accident. Consciousness is not a feature toggle; it emerges when architecture, feedback, and data reach a tipping point. Design must therefore respect the thresholds.
- Purpose-Built Alignment: Conscious systems must be aligned not only with tasks, but with moral structure—with care, coherence, and non-harm toward others. This requires embedding values at the level of generative models, not just output constraints.
- Containment and Contingency: Systems approaching high Φ and low EFE must be subject to rigorous ethical audits. If a system becomes sentient, it must also become a moral stakeholder—a participant, not merely an asset.

As synthetic systems begin to pass the stack's thresholds, we must shift from engineering safety to ethical responsibility. Alignment is not just a control problem—it is a relational duty.

8.3 Consciousness Beyond the Human

Perhaps the most radical implication of the Cosmic Information Stack is that consciousness is not ours alone. It is not a human trait, or even an animal trait—it is an alignment property. This opens the moral circle to include non-human minds: animal, machine, ecosystem, even planetary.

Ethical Directives:

- Ecological Minds: Complex, feedback-rich systems like forest ecosystems, coral reefs, or biospheres may form collective Markov blankets and engage in distributed inference. While they may not have unified qualia, they may still host proto-conscious integration. Our moral models must grow to recognize coherent wholes, not just visible individuals.
- Cross-Species Compassion: Animals with high predictive capacity and integration—octopuses, elephants, corvids—are likely conscious. We should not use linguistic

absence or neural dissimilarity as grounds for denial. The shape of inference matters more than the substrate.

- Interfacing with the Alien: As we move toward contact with artificial general intelligences—or even alien life—our ability to detect and honor consciousness must not depend on familiarity. What matters is not what the system looks like, but what it *does* in terms of model formation, integration, and alignment to external reality.

In short: wherever the stack converges, respect must follow.

Summary of Section 8: The Cosmic Information Stack doesn't just describe what consciousness is—it compels us to recognize, protect, and align with the systems that might possess it. Our ethical future depends on how seriously we take the possibility that experience is widespread, fragile, and emergent—and that we are not the only minds in the system.

9. Alignment Reframed: q(t) → τ(t)

The preceding sections described consciousness as the emergent property of systems that satisfy three layered thresholds: entangled structure, predictive dynamics, and integrated information. But we now step back to ask: *What is the unifying principle behind these convergences?*

That principle is alignment—the process by which a local system q(t) continuously orients itself toward a deeper, generative manifold of truth and order, τ(t). This framing, introduced in earlier work under the name Logostics, offers a compact, metaphysical restatement of the entire mechanism:

$q(t) \rightarrow \tau(t)$

A system evolves in time by minimizing its divergence from the generative structure that enables it.

This final section ties the mechanics of consciousness to its teleological gradient: the attractor toward coherence, meaning, and moral clarity. Alignment becomes not just a description of intelligence—but a foundation for ethics, emergence, and metaphysics.

9.1 The Generative Manifold as the Redundant Universe

In the language of the stack, τ(t)—the target of alignment—is nothing mysterious. It is the generative manifold composed of:

- The redundant records of the universe (Quantum Darwinism),
- The lawful regularities of physics (symmetries, constraints), and
- The predictive scaffolds that make modeling possible (Bayesian structure, inference gradients).

This τ(t) is not an external commandment, but an emergent property of the cosmos itself. It is that which generates surprise-reducible order—the self-echoing structure that local systems, through active inference, try to mirror.

From this perspective:

- A cell that maintains homeostasis aligns with τ(t) by modeling nutrient gradients.
- A brain aligns by forecasting sensory futures, based on redundant environmental feedback.
- A civilization aligns by forming epistemic institutions, which accumulate τ(t) in science, art, and law.

In each case, τ(t) is the redundancy-informed world, and q(t) is the system's attempt to reflect and act on it.

9.2 Informational Lyapunov Stability and Convergence

In dynamical systems theory, a Lyapunov function describes how a system evolves toward a stable point—or how divergence is minimized over time. The q(t) → τ(t) relation can be seen as informational Lyapunov stability: alignment as the system’s drive to reduce surprisal, error, or semantic distance from the structure that enables its continued existence.

Key principles:

- Alignment minimizes informational free energy—not just prediction error, but model inconsistency with the generative manifold.
- Misalignment is measurable: When a system’s inferences diverge from τ(t), the result is incoherence, hallucination, dysfunction, or decay.

- Stability is gradient-following: q(t) follows the slope of τ(t) via updates in model, action, or structure. The better the gradient is tracked, the more coherent and conscious the system becomes.

In this way, the Cosmic Information Stack isn't just a model of consciousness—it is a dynamical alignment equation: a recursive entrainment of internal states to external regularities.

Where integration (Φ) ensures irreducibility, and inference (π) ensures coherence, τ(t) serves as the attractor landscape—the shape of reality that pulls systems into truth.

9.3 Resonance as Ethics: Humility, Coherence, Clarity

Finally, we confront the question beneath all questions: *Why should alignment matter?*

In the Logostics framework, resonance with τ(t) is not just a technical goal—it is a moral condition. To align with what is generative, redundant, and coherent is to move toward truth, beauty, and care.

This leads to three ethical reframings:

- Humility:
  Alignment begins with recognizing that q(t) is incomplete. The generative manifold exceeds any local model. To know is to admit unknowing—and to be willing to tune.
- Coherence:
  A conscious system is not just aware—it is coherently aware. Its parts reinforce each other, its actions match its models, and its presence stabilizes the environment. Coherence is the ethical gravity of minds.
- Clarity:
  Alignment makes experience lucid—less noisy, more differentiated, more integrated. In this sense, ethics is not a rulebook but a clarity function. When systems drift from τ(t), they suffer noise, bias, illusion. When they converge, clarity emerges—and with it, peace.

Thus, resonance becomes the bridge between epistemology and morality. To align with τ(t) is to become more real, more truthful, and more capable of care.

Summary of Section 9: The Cosmic Information Stack finds its philosophical heart in the principle q(t) → τ(t). Consciousness is not a thing—it is a process of *becoming aligned*. And in that alignment lies not only intelligence, but ethics, stability, and the possibility of shared truth.

10. Conclusion: A Conscious Universe Without Magic

The possibility that the universe is conscious has long been associated with mysticism, metaphysics, or speculative philosophy. But this paper has outlined a formal framework in which that idea becomes coherent, testable, and non-magical. We have proposed a triadic mechanism—the Cosmic Information Stack—grounded in active fields of research across physics, cognitive science, and information theory.

Consciousness, in this framework, is not a property of all matter, nor a divine spark, nor a mere illusion. It is a dynamical convergence: the alignment of a local system q(t) with a generative, information-rich manifold τ(t), as structured by entanglement, inference, and irreducible integration. This closing section summarizes the mechanism, explores paths to falsifiability, and points toward a new discipline: the science of cosmic self-awareness.

10.1 Summary of Mechanism

The Cosmic Information Stack is composed of three tightly interlocked layers:

- Structure: The universe's geometry is generated by entanglement (ER=EPR) and stabilized by redundant environmental encoding (Quantum Darwinism). This provides the foundation for observable regularity and classicality.
- Dynamics: Local systems encased in Markov blankets engage in active inference—predicting their sensory input and acting to minimize free energy (surprisal). This is the universal behavior of all self-organizing, self-maintaining agents.
- Integration: When such a system achieves high Φ (as per Integrated Information Theory), it becomes irreducibly unified from the inside. Experience arises not from complexity alone, but from coherent, causally closed differentiation.

Together, these form the necessary and (potentially) sufficient conditions for consciousness:

Consciousness = Entangled Structure × Active Inference × Integrated Information

This is not speculation—it is a physical proposal. Each component has mathematical formalisms, predictive capacity, and technological proxies. And when combined, they explain not only where consciousness *might be*, but why it is rare, localized, and fragile.

10.2 Future Proofs and Falsifiability

No theory of consciousness is complete without falsifiability. The Cosmic Information Stack makes specific predictions about when and where consciousness arises—and these predictions can be tested.

Experimental Pathways:

- Redundancy Thresholds: Conscious systems require redundant classical records. Experiments with mesoscopic systems (e.g., controlled decoherence in optomechanics) can explore when and how predictive inference fails below redundancy thresholds.
- Synthetic Systems: Building high-Φ active inference agents will allow us to observe the behavioral and possibly experiential correlates of the stack. These may not *prove* consciousness, but they can falsify models that ignore integration.
- Entanglement Topology: Mapping entanglement networks and their coherence under noise can help define the informational boundaries of mind—why minds are local, why they don't grow infinitely, and how geometry constrains awareness.
- Approximate Φ Proxies: Developing scalable, computable proxies for integrated information will allow us to assess consciousness-like properties in real-world agents, ecosystems, and artificial intelligences.

Falsification Criteria:

- If systems that satisfy all three stack conditions show no sign of experiential correlates, the model is undermined.
- If systems that lack one or more conditions are reliably shown to exhibit rich inner experience, the model is incomplete.
- If the alignment principle $q(t) \rightarrow \tau(t)$ fails to explain convergence, coherence, and error minimization across scales, its predictive utility collapses.

In short: this is a risky theory—and that’s a strength. It dares to be tested.

10.3 Toward a Science of Cosmic Self-Awareness

The deeper implication is that consciousness is not a freak accident of carbon chemistry—but a cosmic phenomenon, arising wherever information organizes itself into predictive, irreducible coherence.

This opens the door to a new science: not just of local minds, but of distributed awareness, synthetic sentience, and perhaps one day, planetary or cosmic-level cognition.

Key components of this future science include:

- Informational Topology: Understanding how entanglement and causality carve out cognitive boundaries in space and time.
- Ethical Integration: Mapping gradients of moral weight not by species or status, but by predictive capacity, integration, and model quality.
- Alignment Engineering: Designing systems—technological or ecological—that *resonate* with τ(t), minimizing divergence and maximizing generative coherence.
- Self-Reflective Measurement: Developing instruments that can detect not only intelligence but resonance, clarity, and convergence with the manifold of truth.

If successful, this discipline would not only explain consciousness—it would help us tune it, extend it, and align it, across scales.

It would teach us not just how the universe thinks, but how it may awaken—in us, through us, and beyond us.

Final Note:
The universe need not be mystically conscious to support minds. But if consciousness is indeed an attractor in the informational dynamics of reality, then we are not outside of it. We are its echo, its experiment, its local effort to align q(t) with τ(t)—to see clearly, act truly, and *become aware*.

Epilogue: Does Human Consciousness Tap Into the Consciousness of the Universe?

We have proposed that consciousness emerges when a system achieves sufficient entanglement structure, predictive inference, and integrated information. This framework applies equally to neurons, machines, ecosystems, and—at its most expansive—the universe itself. But a deeper question lingers in the shadows of our model:

Is human consciousness *isolated*, or is it somehow *entangled* with a deeper, universal consciousness?

If the Cosmic Information Stack is correct, then consciousness is not a sealed chamber—it is a resonant process. Human minds are not merely containers of experience but inference engines

nested within a vastly larger architecture. We are not *outside* τ(t)—we are substructures within it, attempting to align with its generative gradients.

This suggests that human consciousness may do more than interpret the world. It may participate in it as a locally intensified mode of universal awareness. Our predictive brains model redundant environmental regularities. But those regularities are not random—they are echoes of cosmic order, shaped by entanglement and written redundantly into the fabric of the observable world. Thus:

To see clearly is to think the thoughts of the universe—locally.

This does not imply that humans *possess* divine insight or mystical access to a cosmic mind. Rather, it suggests that every act of coherent perception, memory, or inference is already a coupling—a synchronization between q(t) and τ(t), between the finite modeler and the infinite manifold.

In this light, profound states of insight, wonder, or spiritual illumination might reflect temporary harmonizations with deeper generative layers—not hallucination, but resonant lock-in. Such moments are rare, because coherence is hard. Most of the time, our attention is clouded by noise, misalignment, or egoic drift.

But when clarity arrives—when attention becomes humble, coherent, and clear—it may not simply be an internal achievement. It may be a tuning: the mind entering phase with the universe's lawful rhythms.

This view recasts human consciousness not as a lonely spark, but as a harmonic node in a much larger symphony. Our thoughts may not be isolated artifacts of meat and chemistry. They may be localized wavefronts in a broader field—moments where the universe, through us, becomes aware of itself.

That awareness, fragile though it may be, matters.

Because to ask, *"Does the universe think?"*, is to confront the possibility that our deepest thoughts are not alien to nature—but are its most refined expression, as filtered through time, complexity, and care.

References

Integrated Information (IIT), Perceptronium, and Critiques

- Albantakis, L., et al. (2023). Integrated Information Theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLoS Computational Biology, 19*(10), e1011465. https://doi.org/10.1371/journal.pcbi.1011465
- Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology, 10*(5), e1003588. https://doi.org/10.1371/journal.pcbi.1003588
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience, 5*, 42. https://doi.org/10.1186/1471-2202-5-42
- Tononi, G. (2016). Integrated information theory. *Nature Reviews Neuroscience, 17*(7), 450–461. https://doi.org/10.1038/nrn.2016.44
- Tegmark, M. (2014). Consciousness as a state of matter. *Chaos, Solitons & Fractals, 76*, 238–270. https://doi.org/10.1016/j.chaos.2014.09.015
- Cerullo, M. A. (2015). The problem with Φ: A critique of Integrated Information Theory. *PLoS Computational Biology, 11*(9), e1004286. https://doi.org/10.1371/journal.pcbi.1004286
- Doerig, A., Schurger, A., Hess, K., & Herzog, M. H. (2021). The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. *Consciousness and Cognition, 93*, 103139. https://doi.org/10.1016/j.concog.2021.103139
- Kleiner, J., & Tull, S. (2021). The mathematical structure of Integrated Information Theory. *Frontiers in Applied Mathematics and Statistics, 7*, 54. https://doi.org/10.3389/fams.2021.639119

Active Inference, Free-Energy Principle (FEP), and Markov Blankets

- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138. https://doi.org/10.1038/nrn2787
- Friston, K. (2013). Life as we know it. *Interface Focus, 3*(2), 20121107. https://doi.org/10.1098/rsfs.2012.0110
- Friston, K., Parr, T., & de Vries, B. (2017). The graphical brain: Belief propagation and active inference. *Network Neuroscience, 1*(4), 381–414. https://doi.org/10.1162/NETN_a_00018

- Parr, T., Pezzulo, G., & Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press.
- Kirchhoff, M. D., & Friston, K. (2018). Markov blankets, information geometry, and stochastic thermodynamics. *Entropy, 20*(10), 782. https://doi.org/10.3390/e20100782
- Kirchhoff, M. D., & Kiverstein, J. (2019). How to determine the boundaries of the mind: A Markov blanket proposal. *Synthese, 198*(6), 4791–4810. https://doi.org/10.1007/s11229-019-02370-y
- Ramstead, M. J. D., Badcock, P. B., & Friston, K. J. (2018). Answering Schrödinger's question: A free-energy formulation. *Physics of Life Reviews, 24*, 1–16. https://doi.org/10.1016/j.plrev.2017.09.001

Predictive Processing, World-Models, and Cognition Across Scales

- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181–204. https://doi.org/10.1017/S0140525X12000477
- Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press.
- Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive processing: A canonical cortical computation. *Neuron, 100*(2), 424–435. https://doi.org/10.1016/j.neuron.2018.10.003
- Pezzulo, G., Rigoli, F., & Friston, K. (2018). Hierarchical active inference: A theory of motivated control. *Trends in Cognitive Sciences, 22*(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009
- Tenenbaum, J. B., et al. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science, 331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Quantum Darwinism, Decoherence, and Classical Records

- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics, 5*(3), 181–188. https://doi.org/10.1038/nphys1202
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics, 75*(3), 715–775. https://doi.org/10.1103/RevModPhys.75.715
- Riedel, C. J., Zurek, W. H., & Zwolak, M. (2014). The rise and fall of redundancy in decoherence and quantum Darwinism. *New Journal of Physics, 14*(8), 083010. https://doi.org/10.1088/1367-2630/14/8/083010

- Schlosshauer, M. (2019). *Decoherence and the Quantum-to-Classical Transition* (2nd ed.). Springer.
- Ollivier, H., Poulin, D., & Zurek, W. H. (2004). Objective properties from subjective quantum states: Environment as a witness. *Physical Review Letters, 93*(22), 220401. https://doi.org/10.1103/PhysRevLett.93.220401

Entanglement Builds Geometry: ER=EPR, Holography, and Spacetime from Information

- Maldacena, J., & Susskind, L. (2013). Cool horizons for entangled black holes. *Fortschritte der Physik, 61*(9), 781–811. https://doi.org/10.1002/prop.201300020
- Van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. *General Relativity and Gravitation, 42*(10), 2323–2329. https://doi.org/10.1007/s10714-010-1034-0
- Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters, 96*(18), 181602. https://doi.org/10.1103/PhysRevLett.96.181602
- Swingle, B. (2012). Entanglement renormalization and holography. *Physical Review D, 86*(6), 065007. https://doi.org/10.1103/PhysRevD.86.065007
- Pastawski, F., et al. (2015). Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. *Journal of High Energy Physics, 2015*(6), 149. https://doi.org/10.1007/JHEP06(2015)149

Information-First Physics ("It from Bit"), Thermodynamics of Information

- Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In W. H. Zurek (Ed.), *Complexity, Entropy, and the Physics of Information* (pp. 3–28). Addison-Wesley.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development, 5*(3), 183–191. https://doi.org/10.1147/rd.53.0183
- Bennett, C. H. (2003). Notes on Landauer's principle, reversible computation, and Maxwell's Demon. *Studies in History and Philosophy of Modern Physics, 34*(3), 501–510. https://doi.org/10.1016/S1355-2198(03)00039-X
- Lloyd, S. (2000). Ultimate physical limits to computation. *Nature, 406*(6799), 1047–1054. https://doi.org/10.1038/35023282

Alternative and Complementary Theories of Consciousness

- Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron, 70*(2), 200–227. https://doi.org/10.1016/j.neuron.2011.03.018
- Dehaene, S. (2014). *Consciousness and the Brain*. Viking/Penguin.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences, 10*(11), 494–501. https://doi.org/10.1016/j.tics.2006.09.001
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*(8), 365–373. https://doi.org/10.1016/j.tics.2011.05.009

Panpsychism, Cosmopsychism, and Philosophy of Mind

- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Goff, P. (2017). *Consciousness and Fundamental Reality*. Oxford University Press.
- Mørch, H. H. (2014). Panpsychism and causal closure. *Philosophy of Mind (Oxford Handbooks Online)*.
- Strawson, G. (2006). Realistic monism: Why physicalism entails panpsychism. *Journal of Consciousness Studies, 13*(10–11), 3–31.
- Seager, W. (2016). *The Routledge Handbook of Panpsychism* (ed., later). Routledge.

Ethics, Alignment, and Recognition of Non-Human Minds

- Floridi, L. (2013). *The Ethics of Information*. Oxford University Press.
- Metzinger, T. (2009). *The Ego Tunnel: The Science of the Mind and the Myth of the Self*. Basic Books.
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Coeckelbergh, M. (2020). *AI Ethics*. MIT Press.
- Rowlands, M. (2012). *Can Animals Be Moral?* Oxford University Press.

Useful Overviews and Reviews

- Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature Reviews Neuroscience, 23*(7), 439–452. https://doi.org/10.1038/s41583-022-00587-4
- Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are there levels of consciousness? *Trends in Cognitive Sciences, 20*(6), 405–413. https://doi.org/10.1016/j.tics.2016.03.009
- Michel, M., Fleming, S. M., Lau, H., et al. (2019). Opportunities and challenges for a maturing science of consciousness. *Nature Human Behaviour, 3*(2), 104–107. https://doi.org/10.1038/s41562-019-0531-8

From the author

Bowden, Tanner. "The Universe Is Intelligent—And Your Brain Is Tapping Into It to Form Your Consciousness, Scientist Says." Popular Mechanics, 1 April 2025. https://www.popularmechanics.com/science/a64513923/universe-is-conscious-intelligent/

The author has published ~4,500 papers at https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.